

AFRL-IF-WP-TR-2005-1562

**EVOLVED MULTIREOLUTION
TRANSFORMS FOR OPTIMIZED IMAGE
COMPRESSION AND
RECONSTRUCTION UNDER
QUANTIZATION**



Frank Moore

**University of Alaska Anchorage
Department of Mathematical Sciences
CAS 154, 3211 Providence Drive
Anchorage, AK 99508**

AUGUST 2005

Final Report for 05 May 2005 – 18 August 2005

Approved for public release; distribution is unlimited.

STINFO FINAL REPORT

**INFORMATION DIRECTORATE
AIR FORCE RESEARCH LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7334**

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YY) August 2005		2. REPORT TYPE Final		3. DATES COVERED (From - To) 05/05/2005 – 08/18/2005	
4. TITLE AND SUBTITLE EVOLVED MULTIREOLUTION TRANSFORMS FOR OPTIMIZED IMAGE COMPRESSION AND RECONSTRUCTION UNDER QUANTIZATION				5a. CONTRACT NUMBER FA8650-05-2-0064	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 61102F	
6. AUTHOR(S) Frank Moore				5d. PROJECT NUMBER WOGA	
				5e. TASK NUMBER IF	
				5f. WORK UNIT NUMBER TA	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Alaska Anchorage Department of Mathematical Sciences CAS 154, 3211 Providence Drive Anchorage, AK 99508				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Information Directorate Air Force Research Laboratory Air Force Materiel Command Wright-Patterson AFB, OH 45433-7334				10. SPONSORING/MONITORING AGENCY ACRONYM(S) AFRL-IF-WP	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S) AFRL-IF-WP-TR-2005-1562	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES Report contains color. PAO Case Number: AFRL/WS-04-0658, 01 Sep 2004.					
14. ABSTRACT State-of-the-art image compression and reconstruction techniques utilize wavelets. Beginning in 2004, however, ongoing research at Wright-Patterson Air Force Base (WPAFB), the University of Alaska Anchorage (UAA), and the Air Force Institute of Technology (AFIT) has demonstrated that a genetic algorithm (GA) is capable of evolving nonwavelet transforms that consistently outperform wavelets when applied to a broad class of images under conditions subject to quantization error. This report describes recent research that builds upon those previous results in each of the following ways: <ol style="list-style-type: none"> 1. First, this research demonstrates that a GA can evolve a single set of coefficients describing a single matched forward and inverse transform pair that can be used at each level of a multiresolution transform to simultaneously minimize the size of the compressed file and the squared error (SE) in the reconstructed file. 2. Second, this research examines the relationship between the specified quantization level and the performance of the evolved transform relative to the wavelet. 3. Third, this research extends the GA to simultaneously evolve k sets of coefficients—one set for each level of a k-level multiresolution transform—that further reduce error in reconstructed images. 4. Fourth, this research attempts to evolve k-level multiresolution transforms using highly specific images as the training population, in hope that the resulting transform would exhibit better performance (in terms of reduced error) on similar images, but poorer performance against dissimilar images. 					
15. SUBJECT TERMS Genetic algorithms, discrete wavelet algorithms, image compression, quantization noise					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT: SAR	18. NUMBER OF PAGES 30	19a. NAME OF RESPONSIBLE PERSON (Monitor) Pat Marshall 19b. TELEPHONE NUMBER (Include Area Code) (937) 255 6548, ext. 3609
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			

ACKNOWLEDGEMENTS

The research described in this paper was performed as part of the Air Force Research Laboratory's 2005 Visiting Faculty Research Program. The author especially recognizes Pat Marshall (AFRL/IFTA) at WPAFB, Dayton, OH, and Barry McKinney (AFRL/IF) at Rome Research Site (RRS), Rome, NY, for their continuing support. Michael Peterson (Wright State University, Dayton, OH), Gary Lamont (AFIT), and Eric Balster (AFRL/IFTA) also contributed to this research.

1. SUMMARY

The purpose of this Visiting Faculty Research Program (VFRP) project was to establish a methodology for using a genetic algorithm (GA) to evolve coefficient sets representing matched forward and inverse multiresolution transform pairs capable of outperforming wavelets for image compression and reconstruction applications subject to quantization error.

First, the GA was extended to evolve a *single* set of coefficients which could be used at each level of a multiresolution analysis (MRA) transform. The resulting transform outperformed the Daubechies-4 (D4) wavelet in terms of reducing aggregate squared error in compressed, reconstructed images at various quantization levels.

Second, the GA was tested using quantization levels of 16, 32, and 64. The results of these tests demonstrate that the performance improvement of evolved transforms over wavelets increases in proportion to the amount of quantization, with the largest improvement occurring for a quantization step of 64.

Third, the GA was used to simultaneously evolve *multiple* sets of coefficients—one set for *each level* of an MRA transform. The resulting transforms outperformed evolved MRA transforms described by a single set of coefficients, and substantially outperformed the D4 wavelet, when subsequently tested against a set of standard test images.

Fourth, the multicoefficient set MRA transform approach was used to train a transform on a highly specific image (in this case, the image of a grey hammer), in hope that the resulting transform would perform very well when tested against similar images (i.e., other hammers), but poorly when tested against dissimilar images (e.g., dollar signs, deoxyribonucleic acid (DNA) models, and eyedroppers). However, these test results suggest that the content of an image has little to do with whether an evolved transform performs well or poorly when compressing and reconstructing images under conditions subject to quantization.

2. INTRODUCTION

Since the late 1980s, engineers, scientists, and mathematicians have used wavelets (Daubechies 1992) to solve a wide variety of difficult problems, including fingerprint compression (Bradley, Brislawn, and Hopper 1993), signal denoising (Donoho 1993), and medical image processing (Aldroubi and Unser 1996). Recent adoption of the Joint Photographic Experts Group's JPEG2000 standard (Taubman and Marcellin 2002) has established wavelets as the principal methodology for image compression and reconstruction. JPEG2000 utilizes wavelets to improve upon the compression capabilities of previous JPEG (International Organization for Standardization (ISO) 1994) and JPEG-LS (ISO 1999) standards.

Wavelets may be described by four sets of coefficients:

1. $h1$ is the set of wavelet numbers for the (forward) discrete wavelet transform (DWT).
2. $g1$ is the set of scaling numbers for the DWT.
3. $h2$ is the set of wavelet numbers for the inverse DWT (DWT^{-1}).
4. $g2$ is the set of scaling numbers for the DWT^{-1} .

For the D4 wavelet, these sets consist of the following floating-point coefficients:

$$\begin{aligned}h1 &= \{-0.129409523, 0.224143868, 0.836516304, 0.482962913\} \\g1 &= \{-0.482962913, 0.836516304, -0.224143868, -0.129409523\} \\h2 &= \{0.482962913, 0.836516304, 0.224143868, -0.129409523\} \\g2 &= \{-0.129409523, -0.224143868, 0.836516304, -0.482962913\}\end{aligned}$$

A two-dimensional (2-D) DWT of a discrete input image \mathbf{f} with M rows and N columns is computed by first applying the one-dimensional (1-D) subband transform defined by the coefficients from sets $h1$ and $g1$ to the columns of \mathbf{f} , and then applying the same transform to the rows of the resulting signal (Taubman and Marcellin 2002, p. 428). Similarly, a 2-D DWT^{-1} is performed by applying the 1-D DWT^{-1} defined by sets $h2$ and $g2$ first to the rows and then to the columns of a previously compressed signal.

A one-level DWT decomposes \mathbf{f} into $M/2$ -by- $N/2$ subimages \mathbf{h}^1 , \mathbf{d}^1 , \mathbf{a}^1 , and \mathbf{v}^1 , where \mathbf{a}^1 is the trend subimage of \mathbf{f} and \mathbf{h}^1 , \mathbf{d}^1 , and \mathbf{v}^1 are its first horizontal, diagonal, and vertical fluctuation subimages, respectively (Walker 1999). Using the MRA scheme (Mallat 1989), a one-level DWT may be repeated $k \leq \log_2(\min(M, N))$ times. Note that the size of the trend signal \mathbf{a}^i at level i of decomposition will be $1/4^i$ times the size of the original image \mathbf{f} (e.g., a three-level transform produces a trend subimage \mathbf{a}^3 that is $1/64^{\text{th}}$ the size of \mathbf{f}). Nevertheless, the trend subimage will typically be much larger than any of the fluctuation subimages; for this reason, the MRA scheme computes a k -level DWT by recursively applying a one-level DWT to the rows and columns of the discrete trend signal \mathbf{a}^{k-1} . Similarly, a one-level DWT^{-1} is applied k times to reconstruct an approximation of the original M -by- N signal \mathbf{f} .

Quantization (Mallat 1998) is the most common source of distortion in lossy image compression systems. Quantization refers to the process of mapping each of the possible

values of a given sampled signal y onto a smaller range of values $Q(y)$. The resulting reduction in the precision of data allows a quantized signal q to be much more easily compressed. The corresponding dequantization step, $Q^{-1}(q)$, produces signal \hat{y} that differs from the original signal y according to a distortion measure ρ . A variety of techniques may be used to quantify distortion; however, if we assume that quantization errors are uncorrelated, then the aggregate distortion in the dequantized signal, $\rho(y, \hat{y})$, may be computed as a linear combination of squared error (SE) for each sample.

3. THE GA

The goal of any effective image compression and reconstruction system is to simultaneously minimize two parameters:

1. The number of bits needed to represent the compressed image produced by the forward transform (i.e., the compressed file size FS).
2. The distortion observed in the reconstructed image produced by the corresponding inverse transform (i.e., the SE).

The purpose of the research described by this paper was to determine whether a GA (Goldberg 1989) could be used to evolve coefficient sets representing nonwavelet transforms capable of outperforming wavelet-based multiresolution transforms under conditions subject to quantization error.

The following parameters characterize the GA developed to achieve this goal:

1. The maximum number of evolved generations $G = 2000$.
2. The size of the evolving population $M = 2000$.
3. The number of multiresolution levels $MR = 3$.
4. The probability of crossover $p_c = 90$ percent.
5. The probability of mutation p_m for any candidate solution was initialized to a user-specified minimum. If the current generation failed to identify a new best-of-run solution, p_m was increased by a selected increment up to a user-specified maximum mutation rate. The training runs described in this paper used $\min(p_m) = 2$ percent, $\max(p_m) = 20$ percent, and a 2 percent increment.
6. The GA trained each transform using a representative 128- by 128-pixel subimage of the standard 512- by 512-pixel “couple.bmp” image. This subimage was chosen based on the results of previous investigations which demonstrated that the resulting evolved transforms generalized well for other images in the test set.

For each of the tests described in this paper, each candidate solution specified the floating-point coefficients for sets $g1$, $h1$, $h2$, and $g2$. The GA seeded the initial population (generation 0) with one exact copy and $M-1$ randomly mutated copies of the D4 wavelet (Ramsey and Grefenstette 1993). Thus, sets $g1$, $h1$, $h2$, and $g2$ of every individual in the population each contained precisely four coefficients. After fitness evaluation, the individual with the best fitness value was copied into position 0 of the next generation (De Jong 1975), while the remaining $M-1$ positions were populated using tournaments of a user-specified number of randomly selected individuals from the current generation (Miller and Goldberg 1995).

Next, the GA performed single-point crossover on adjacent pairs of individuals with probability p_c . The crossover operator randomly selected one of the four coefficient sets, and then randomly selected a crossover point within that set. The coefficients appearing at or below the selected crossover point in the selected coefficient set from each parent individual were exchanged to create two new candidate solutions.

Finally, mutation was performed on each individual with probability p_m . For this investigation, mutation consisted of multiplying a randomly selected coefficient from a randomly selected set ($g1$, $h1$, $h2$, or $g2$) by a factor randomly selected from a Gaussian distribution between 0.0 and 2.0 and centered upon 1.0. Previous studies suggested that an occasional sign change of coefficients could also be beneficial in reducing SE; therefore, with a (typically very small) user-specified probability, the mutation operator used in this study also negated the mutated coefficient.

4. FITNESS

This study utilized two key quantities to measure fitness:

1. File Size Ratio (FSR) = (the size of the file compressed by the evolved forward transform) / (the size of the file compressed by the DWT)
2. Error Ratio (ER) = (the SE in the image reconstructed by the evolved inverse transform) / (the SE in the image reconstructed by the DWT⁻¹).

Previous research (Babb, Becke, and Moore 2005) established the existence of a nearly linear Pareto-optimal front (Goldberg, Horn, and Nafpliotis 1994) describing the tradeoff between these two conflicting criteria. For this study, the fitness of each candidate solution against a particular image from the training set was measured as follows:

1. First, the GA used the forward transform coefficients specified by the candidate solution to compress the image.
2. Next, compressed image was quantized using the quantization step defined for the current training run, encoded, decoded, and dequantized.
3. Finally, the GA reconstructed the image using the inverse transform coefficients specified by the candidate solution, and calculated the FSR and ER.

Given a training population consisting of one or more images, this study used the following algorithm to estimate the fitness of a given candidate solution:

```
fitness = 0;
for each image in the training population
  if (FSR > 1.0 && ER > 1.0) fitness += FSRA + ERB; // case 1
  else if (FSR > 1.0 && ER ≤ 1.0) fitness += FSRC + ER; // case 2
  else if (FSR ≤ 1.0 && ER > 1.0) fitness += FSR + ERD; // case 3
  else fitness += FSRE + ERF; // case 4
```

Here, A, B, C, D, E, and F are user-specified constants greater than 1.0. (For this study, A = B = C = D = 8 and E = F = 16.) Lower fitness values are better. Cases 1 and 2 thus explicitly penalize the fitness of evolved forward and inverse transform pairs that increase the size of the compressed file; similarly, cases 1 and 3 penalize transforms that result in higher SE. On the other hand, case 4 explicitly rewards evolved transforms that simultaneously reduce both compressed file size and SE, relative to the wavelet.

5. ONE TRANSFORM FOR ALL MRA LEVELS

Previous research focused upon evolving coefficients for either an inverse nonwavelet transform [(Moore, Marshall, and Balster 2005), (Moore 2005)] or a matched forward and inverse nonwavelet transform pair (Babb, Becke, and Moore 2005) that reduced mean squared error (MSE) relative to the performance of a standard wavelet DWT^{-1} or matched DWT/DWT^{-1} pair applied to the same images under conditions subject to a quantization step of 64. The resulting evolved transforms consistently reduced MSE by as much as 25 percent when applied to images from both the training and test sets.

Unfortunately, none of these previous studies involved MRA; instead, coefficients were optimized only for single-level image decomposition and/or reconstruction transforms. Subsequent testing demonstrated that the performance of transforms evolved at a single level of resolution degraded substantially when subsequently tested in a multiresolution environment.

In practice, virtually all wavelet-based compression schemes entail several stages of decomposition. Typical wavelet-based MRA applications compress a given image by recursively applying the $h1$ and $g1$ coefficients a defining single DWT at each of k levels. Image reconstruction requires k recursive applications of the $h2$ and $g2$ coefficients defining the corresponding DWT^{-1} . The JPEG2000 standard allows between $0 \leq k \leq 32$ DWT stages; near-optimal performance on full-resolution images is reported for $D = 5$ levels (Taubman and Marcellin 2002, p. 429).

The first goal of this research effort was to determine whether a GA could evolve a single set of coefficients for a matched evolved forward and inverse transform pair satisfying each of the following conditions:

1. The evolved coefficients were intended for use at each and every level of decomposition by a matched multilevel transform pair.
2. The evolved forward transform produced compressed files whose size was less than or equal to those produced by the DWT.
3. When applied to the compressed file produced by the matching evolved forward transform, the evolved inverse transform produced reconstructed images whose SE was less than or equal to the SE observed in images reconstructed by the DWT^{-1} from files previously compressed by the DWT.

Previous research also failed to establish the relationship between the specified quantization level and the performance of evolved transforms relative to that of wavelets. For this reason, the second goal of this research was to establish the performance enhancement of evolved transforms over wavelets as a function of quantization.

5.1 Test Results

To achieve the first two research goals described above, three training runs were performed. These runs differed only according to the specified quantization level. Test results (Table 1) confirmed the GA's ability to evolve coefficients for a single transform that exhibited optimized performance when applied to every level of a multiresolution

transform. For Test 3, the GA evolved coefficients that simultaneously reduced SE by almost 6.5 percent while maintaining a compressed file size smaller than that produced by the D4 wavelet.

Table 1. Improvement of Evolved Transforms over Wavelets as a Function of Quantization Level

<u>Test</u>	<u>Q</u>	<u>File Size / SE (DWT)</u>	<u>File Size / SE (evolved)</u>	<u>Improvement (SE)</u>
1	16	2162 / 447535.38	2161 / 438555.80	-2.006 percent
2	32	1229 / 1093462.63	1228 / 1047424.95	-4.210 percent
3	64	667 / 2527851.95	666 / 2364332.55	-6.469 percent

These results, combined with similar observations from previous studies, e.g., Babb, Becke, and Moore 2005, substantiate each of the following claims:

1. A GA is capable of evolving matched forward and inverse transform pairs that outperform wavelets at a specified quantization level.
2. The performance improvement of evolved transforms over wavelets increases in proportion to the level of quantization.

Note that, although coefficient sets $g1$, $h1$, $h2$, or $g2$ for every candidate solution were initialized to randomly perturbed copies of the coefficients defining the D4 wavelet, 45 of the 48 coefficients (93.75 percent) have undergone some change during the evolutionary process. This result corroborates previous test data and underscores the fact that the search space immediately adjacent to the D4 wavelet appears to be rich with nonwavelet transforms that may outperform wavelets under conditions subject to quantization error. Close inspection of these coefficients reveals an interesting phenomenon: in general, the greater the amount of quantization, the greater the difference between evolved coefficients and wavelet coefficients. Also interesting is the fact that none of the evolved coefficients differed in sign from the corresponding wavelet coefficient. Whatever benefits the sign change mutation may have had during previous studies (without multiresolution) appears to have been eliminated during the evolution of a single set of coefficients for the optimized multiresolution transforms identified during this study.

Table 2 tabulates the coefficients produced by the training runs from Table 1 and notes the percentage change in each evolved coefficient from sets $g1$, $h1$, $h2$, and $g2$ relative to the corresponding coefficient from the D4 wavelet.

**Table 2. Evolved Coefficients and Percentage Change from D4 Coefficients:
One Transform for All MRA Levels**

<u>Test</u>	<u>Set</u>	<u>Coefficients (Percentage magnitude difference from D4 coefficients)</u>		
1	<i>g1</i>	-0.4831928406	(+0.05 percent)	
		0.8365163040	(unchanged)	
		-0.2277694276	(+1.62 percent)	
		-0.1289164106	(-0.38 percent)	
		<i>h1</i>	-0.1294917987	(+0.06 percent)
			0.2242505778	(+0.05 percent)
	0.8398953785		(+0.40 percent)	
	<i>h2</i>	0.4793849332	(-0.74 percent)	
		0.4830777810	(+0.02 percent)	
		0.8291240048	(-0.88 percent)	
	<i>g2</i>	0.2251359248	(+0.44 percent)	
		-0.1227483711	(-5.15 percent)	
		-0.1318678078	(+1.90 percent)	
		-0.1988169414	(-11.30 percent)	
		0.8344765791	(-0.24 percent)	
		-0.4649087239	(-3.74 percent)	
	2	<i>g1</i>	-0.4851359202	(+0.45 percent)
			0.8394985463	(+3.57 percent)
-0.2269758897			(+1.26 percent)	
-0.1264251009			(-2.31 percent)	
<i>h1</i>			-0.1300256428	(+0.48 percent)
			0.2240904941	(-0.02 percent)
		0.8398953785	(+0.40 percent)	
<i>h2</i>		0.4798481072	(-0.64 percent)	
		0.4845747470	(+0.33 percent)	
		0.8203178205	(-1.94 percent)	
<i>g2</i>		0.2232898873	(-0.38 percent)	
		-0.1133667585	(-12.40 percent)	
		-0.1312233947	(+1.40 percent)	
		-0.1681967819	(-24.96 percent)	
		0.8352313868	(-0.15 percent)	
		-0.4547615370	(-5.84 percent)	
3		<i>g1</i>	-0.5008454816	(+3.70 percent)
			0.8365163040	(unchanged)
	-0.2158388997		(-3.71 percent)	
	-0.1314604618		(+1.58 percent)	
	<i>h1</i>		-0.1285400096	(-0.67 percent)
			0.2241438680	(unchanged)
		0.8377104749	(+0.14 percent)	
	<i>h2</i>	0.4827317796	(-0.05 percent)	
		0.4896825540	(+1.39 percent)	
		0.8082258125	(-3.38 percent)	
	<i>g2</i>	0.2183220074	(-2.60 percent)	
		-0.1034099818	(-20.09 percent)	
		-0.1443190513	(+11.52 percent)	
		-0.1399062106	(-37.58 percent)	
		0.8240345243	(-1.49 percent)	
		-0.4365732803	(-9.61 percent)	

6. DIFFERENT TRANSFORMS FOR EACH OF K MRA LEVELS

Traditionally, MRA techniques apply the same DWT at every level of decomposition. As discussed above, the trend subimage \mathbf{a}^k produced by the forward transform at level k becomes the input signal to the forward transform at level $k+1$. Since subimage \mathbf{a}^{k+1} generally corresponds to the vertically low-frequency and horizontally low-frequency (LL) subband output produced by the application of a separable low-pass filter to \mathbf{a}^k , the frequency content of \mathbf{a}^{k+1} will differ substantially from that of \mathbf{a}^k . Since the performance of any particular transform also varies with frequency, it seemed plausible that a multiresolution transform could be enhanced by evolving a different transform at each resolution level. Each evolved transform would, in effect, be targeted towards a particular frequency range. In light of these observations, the third goal of this research was to use the GA to evolve a different set of optimized $g1$, $h1$, $h2$, and $g2$ coefficients for each resolution level of a matched forward and inverse MRA transform.

6.1 Test Results

Table 3 summarizes the result of a preliminary test, which used a quantization step of 64. This result suggests that additional reduction in aggregate squared error may be obtained by evolving matched forward and inverse transform coefficients that are optimized to perform at a specified MRA level.

**Table 3. Improvement of Evolved Transforms over Wavelets:
One Coefficient Set Per Multiresolution Level**

<u>File Size / SE (DWT)</u>	<u>File Size / SE (evolved)</u>	<u>Improvement (SE)</u>
667 / 2527851.95	666 / 2285856.38	-9.573 percent

Table 4 lists the $g1$, $h1$, $h2$, and $g2$ coefficients evolved at each level. Note that a variety of conventions for designating multiresolution levels appear in the literature. Table 4 labels the first-level forward transform, i.e., the transform applied to the original image, “MR Level 1,” the second-level forward transform (applied to the first trend subsignal \mathbf{a}^1) “MR Level 2,” and the third-level forward transform (applied to the second trend subsignal \mathbf{a}^2) “MR Level 3.” Inverse transform coefficients are applied in reverse order.

GAs have proven especially useful for searching unpredictable search spaces. The result tabulated in Table 4 illustrates the utility of using a GA to search the transform space immediately adjacent to the wavelet for nonwavelet transforms exhibiting optimized performance. Interesting aspects of the evolved coefficients include the following:

1. The solution space for this problem is enormous, requiring simultaneous optimization of 48 floating-point coefficient values. Only 36 coefficients (75 percent) had undergone any modification in the best-of-run solution (i.e., the remaining 12 coefficients were equal to the corresponding coefficients from the D4 wavelet). This result suggests that a much larger run of the GA will be necessary to encourage further exploration of this solution space.

**Table 4. Evolved Coefficients and Percentage Change from D4 Coefficients:
One Coefficient Set Per Multiresolution Level**

<u>MR Level</u>	<u>Set</u>	<u>Coefficients (Percentage magnitude difference from D4 coefficients)</u>	
1	<i>g1</i>	-0.5046268794	(+4.49 percent)
		0.8088144645	(-3.31 percent)
		-0.1943402200	(-13.30 percent)
	<i>h1</i>	-0.0914463013	(-29.34 percent)
		-0.1294095230	(unchanged)
		0.2241438680	(unchanged)
		0.8365163040	(unchanged)
		0.4829629130	(unchanged)
	<i>h2</i>	0.4703775944	(-2.61 percent)
		0.8188892589	(-2.11 percent)
		0.2425934518	(+8.23 percent)
		-0.1066458058	(-17.59 percent)
	<i>g2</i>	-0.1037258819	(-19.85 percent)
		-0.0636110754	(-71.62 percent)
		0.7472791434	(-10.67 percent)
-0.2806335741		(-58.11 percent)	
2	<i>g1</i>	-0.5492671638	(+13.73 percent)
		0.8858099497	(+5.89 percent)
		-0.2247833668	(+2.85 percent)
		-0.1272344181	(-1.59 percent)
	<i>h1</i>	-0.1237069434	(-4.41 percent)
		0.2251024349	(0.43 percent)
		0.8365163040	(unchanged)
		0.4833077634	(+0.07 percent)
		0.4977626976	(+3.06 percent)
	<i>h2</i>	0.7970895235	(-4.71 percent)
		0.2136541208	(-4.68 percent)
		-0.0881120885	(-68.09 percent)
		-0.1466747685	(+13.34 percent)
	<i>g2</i>	-0.1045140657	(-53.37 percent)
		0.7647935800	(-8.57 percent)
-0.4169790429		(-13.66 percent)	
-0.4829629130		(unchanged)	
3	<i>g1</i>	0.8365163040	(unchanged)
		-0.2241438680	(unchanged)
		-0.1366008437	(-1.59 percent)
		-0.1340944162	(+3.62 percent)
	<i>h1</i>	0.2241438680	(unchanged)
		0.8365163040	(unchanged)
		0.4829629130	(unchanged)
		0.4820421058	(-0.19 percent)
		0.8189838421	(-2.10 percent)
	<i>h2</i>	0.2169559598	(-3.21 percent)
		-0.1233530196	(-4.68 percent)
		-0.1294095230	(unchanged)
		-0.1866018343	(-16.75 percent)
	<i>g2</i>	-0.8181892891	(-2.19 percent, <u>change of sign</u>)
		-0.5041215305	(+4.38 percent)

2. All but one of the 48 coefficients (97.92 percent)—the third coefficient of g_2 for the MR Level 3 inverse transform from Table 4—retain the sign of the corresponding D4 wavelet coefficients.
3. The h_1 coefficients for all three multiresolution levels remain virtually unchanged.
4. The g_1 coefficients vary wildly from one multiresolution level to the next. While g_1 coefficients at MR Level 3 remain virtually unchanged, the magnitude of the third and fourth g_1 coefficients at MR Level 1 have substantially shrunk, and the magnitude of the first g_1 coefficient at MR Level 2 has substantially grown.
5. Variance in the g_2 coefficients is even greater than that of the g_1 coefficients: the magnitude of the second and fourth g_2 coefficients at MR Level 1 decreased by more than 71 percent and 58 percent, respectively, and the second g_2 coefficient at MS Level 2 decreased in magnitude by over 53 percent. While all four g_2 coefficients at MR Level 1 substantially decreased in magnitude, the first g_2 coefficient at MR Level 2 actually increased by more than 13 percent. In addition, changes to g_2 coefficients at MR Level 3 appear to be almost chaotic.

7. EVOLUTION OF GA TRANSFORMS TO DETECT SPECIFIC SUBIMAGES

Previous research suggested that it may be possible use a GA to evolve a matched DWT/DWT^{-1} pair capable of highlighting a specific subimage within a larger scene. Such a transform might be specifically designed to highlight, for example, each existence of a particular type of vehicle in a series of satellite images. To begin to determine whether an evolved transform might be used to solve this problem, a large set of thumbnail-style images was downloaded from the internet. Several of these images are shown in Figure 1. In addition, several subimages extracted from images commonly found in the wavelet literature were used. Some of these subimages are shown in Figure 2, where most and least indicate the subimage with the highest and lowest energy, respectively.

Since HammerGrey and HammerRed are quite similar, and are distinct from each of the other images, HammerGrey was used to train the transform coefficients. As with the previous tests, $MR = 3$ multiresolution levels were used with a quantization step of 64, and the initial population was seeded with mutated copies of the D4 wavelet. After



Figure 1. Highly Specific Thumbnail-style Images: 48- by 48-pixel Bitmap Files (.bmp)

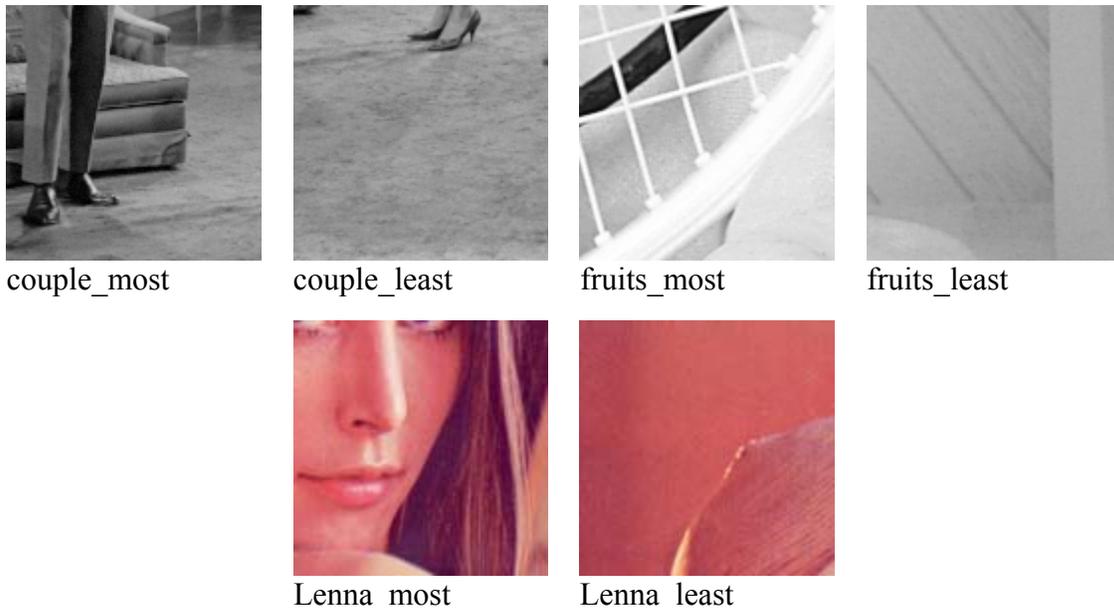


Figure 2. 128- by 128-pixel Subimages

evolving $M = 1000$ candidate solutions over $G = 2000$ generations, the GA produced the best-of-run coefficients shown in Table 5.

Table 5. Coefficients Evolved Using HammerGrey

<u>MR Level</u>	<u>Set</u>	<u>Coefficients</u>	
1	<i>g1</i>	-0.5015777213	
		0.8408070988	
		-0.1987352786	
	<i>h1</i>	-0.1361857608	
		-0.1380635886	
		0.2058384465	
	<i>h2</i>	0.8365163040	
		0.4919633499	
		0.4878009237	
	<i>g2</i>	0.8359185788	
		0.2221604921	
		-0.1295116018	
	2	<i>g1</i>	-0.0890826154
			-0.1135785638
			0.8145441704
<i>h1</i>		-0.4094149435	
		-0.4829629130	
		0.8365163040	
<i>h2</i>		-0.2230739628	
		-0.1294095230	
		-0.1294095230	
<i>g2</i>		0.2267955049	
		0.8365163040	
		0.4829629130	
<i>h2</i>		0.4911973331	
		0.8257232257	
		0.2199939874	
<i>g2</i>	-0.1181444442		
	-0.1230087228		
	-0.1359064655		
3	<i>g1</i>	0.8292301190	
		-0.4342508398	
		-0.4829629130	
	<i>h1</i>	0.8365163040	
		-0.2226445611	
		-0.1259429845	
	<i>h2</i>	-0.1281088918	
		0.2342685134	
		0.8381081538	
	<i>g2</i>	0.4822023422	
		0.4829629130	
		0.8343225514	
	<i>h2</i>	0.2233419196	
		-0.1282597086	
		-0.1577490928	
<i>g2</i>	-0.1950099635		
	0.8275370985		
	-0.4647737675		

Next, these coefficients were applied to the images and subimages shown in Table 5. Table 6 tabulates the file size and SE for the evolved transform and the D4 transform, and notes the percentage change in each.

Table 6. Performance of Coefficients Evolved Against HammerGrey when Tested Against Other Images

<u>Figure</u>	<u>File Size/SE (DWT)</u>	<u>File Size/SE (evolved)</u>	<u>Improvement (SE)</u>
HammerGrey	241 / 624730.12	240 / 569595.39	-8.825 percent
HammerRed	183 / 663310.12	184 / 640356.89	-3.460 percent
Books	483 / 1390262.01	483 / 1422944.97	+2.351 percent
DNA	328 / 1248011.21	331 / 1219049.78	-2.321 percent
Dollar	76 / 134424.91	77 / 128736.47	-4.232 percent
Eyedropper	134 / 392939.11	123 / 378155.99	-3.762 percent
Fish	136 / 447656.67	137 / 429758.13	-3.998 percent
Floppy	216 / 507150.66	221 / 517934.93	+2.216 percent
Couple_most	667 / 2527851.95	666 / 2478155.70	-1.966 percent
Couple_least	424 / 1302731.91	428 / 1283673.14	-1.463 percent
Fruits_most	715 / 2062588.47	707 / 2084257.73	+1.051 percent
Fruits_least	346 / 407252.59	346 / 414928.81	+1.885 percent
Lenna_most	508 / 1588478.43	507 / 1533484.44	-2.203 percent
Lenna_least	341 / 690361.71	341 / 692251.41	+0.274 percent

These results indicate that an MRA transform trained against a single distinct image (HammerGrey) averages only 1.202 percent SE reduction over the entire set of test images. In spite of achieving nearly 9 percent SE reduction on the training image, the evolved transform only achieved a 3.460 percent reduction on a nearly identical test image (HammerRed), and performed equally well on such dissimilar test images as Dollar (4.232 percent reduction), Eyedropper (3.762 percent reduction), and Fish (3.998 percent reduction). In short, none of these results lend much credibility to the initial hypothesis. The fact that the performance of this evolved transform approximately equaled that of the standard D4 wavelet on similar and dissimilar images appears to indicate that the goal of evolving a transform to highlight a particular subimage within a larger scene cannot be achieved using the GA-based approach described in this paper.

8. CONCLUSIONS

This research demonstrated each of the following key points:

1. A GA could evolve coefficients describing a single matched forward and inverse transform pair that was capable of outperforming a similarly structured standard DWT for a specified MRA level.
2. The advantage of using evolved MRA transforms over DWTs increased in proportion to the specified quantization level.
3. The GA could be extended to evolve a different set of coefficients for each level of an evolved MRA transform. At a quantization level of 64, transforms described by multiple coefficient sets outperformed transforms that applied a single set of evolved coefficients at each level of MRA analysis, and substantially outperformed the D4 wavelet.
4. Considerable additional testing will be necessary over a variety of training scenarios to determine whether any discernable pattern in the evolved coefficients emerges.
5. Training on a visually specific image does not result in the evolution of transforms suitable for finding similar images among a larger image set.

Statistical validation of the preliminary results described above will necessitate completion of a much greater number of training runs. Such tests will necessitate several weeks of computation. The relatively short duration of the Air Force Research Laboratory's Summer 2005 VFRP made completion of these tests infeasible.

9. FUTURE DIRECTIONS

During the course of this investigation, it became clear that the amount of computation needed to establish an upper bound on the performance enhancement to be gained via evolved transforms far exceeded available resources. The results summarized above should be interpreted as having demonstrated the feasibility of using GAs to evolve optimized MRA transforms. Close inspection of training run results indicate that the GA were continuing to make evolutionary progress, even as the number of generations approached the selected maximum number (G). It thus appears likely that the performance of evolved MRA transforms relative to wavelets could be further enhanced, merely by increasing the size of each training run.

Similarly, the results of this investigation place no upper bound on the performance enhancement to be gained from evolving multiple sets of $g1$, $h1$, $h2$, and $g2$ coefficients for a multiresolution transform (one set per multiresolution level). The results of this investigation merely demonstrated the benefits of this approach over applying a single set of (standard wavelet or evolved) coefficients at every level of decomposition. Larger scale runs may be able to evolve multiple sets of coefficients for MRA transforms that result in considerably greater SE reduction for various classes of images.

Clearly, future research must determine ways to begin to converge upon an upper bound of performance enhancement. The obvious first step towards this goal would be to exploit available clusters (Bonham and Parmee 1999) and/or supercomputers (Brinkman, Merkle, Lamont, and Pachter 1993) to greatly accelerate GA computation. Other steps should employ various benchmarking techniques to determine whether evolutionary progress can be improved by using smaller populations (Monsieurs and Flerackers 2003), smaller training images, different crossover and mutation schedules, different methods of creating training populations, etc. In addition, a self-tuning GA (Galaviz-Casas and Kuri 1996) might be able to adjust control parameters more effectively during the course of a training run. Finally, the use of alternative numerical optimization techniques (Corne, Dorigo, and Glover 1999), including differential evolution (Price, Storn, and Lampinen 2005), should be considered for future investigations.

It is possible that the GA methodology established by this research could be used to evolve post-reconstruction transforms that do a much better job of approximating the original image than current evolved inverse transforms. Such transforms may take advantage of fitness measures, such as the Statistical, Sampling Inventory Method (SSIM) metric (Wang et al. 2004), that more accurately model human visual system response.

The research suggested that *visual content* of images has no discernable impact the performance of evolved transforms, as measured by the SE metric: transforms evolved against the hammerGrey image performed equivalently to the D4 wavelet on such diverse images as hammerRed, eyedropper, and dollarSign. However, prior research also suggested that the performance of transforms trained against signals with similar *energy distribution* will deteriorate when tested against signals having a very different energy

distribution. Future research should focus upon the problem of evolving transforms that are sensitive to specific energy distributions. This line of research may ultimately lay the foundation for creating transforms capable of highlighting subimages having specific energy distributions within larger scenes.

10. REFERENCES

- Aldroubi, A. and M. Unser (eds.) 1996. *Wavelets in Medicine and Biology*, CRC Press.
- Babb, B., S. Becke, and F. Moore 2005. Evolving Optimized Matched Forward and Inverse Transform Pairs via Genetic Algorithms, *Proceedings of the 48th IEEE International Midwest Symposium on Circuits and Systems: Cincinnati, OH, August 7-10, 2005*, IEEE Circuits and Systems Society. In press.
- Bonham, C. and I. Parmee 1999. An Investigation of Exploration and Exploitation within Cluster Oriented Genetic Algorithms (COGAs), *Proceedings of the Genetic and Evolutionary Computation Conference*, 2: 1491-1497, Morgan Kaufmann.
- Bradley, J., C. Brislawn, and T. Hopper 1993. The FBI Wavelet/Scalar Quantization Standard for Gray-Scale Fingerprint Image Compression, *SPIE Vol. 1961: Visual Information Processing II (1993)*: 293-304, SPIE.
- Brinkman, D., L. Merkle, G. Lamont, and R. Pachter 1993. Parallel Genetic Algorithms and Their Application to the Protein Problem, *Intel Supercomputer User Group 1993 Conference (ISUG)*, 281-293, Intel Scientific Computers.
- Corne, D., M. Dorigo, and F. Glover 1999. *New Ideas in Optimization*, ISBN 0077095065, McGraw-Hill.
- Daubechies, I. 1992. *Ten Lectures on Wavelets*, SIAM.
- Davis, G. and A. Nosratinia 1998. Wavelet-Based Image Coding: An Overview, *Applied and Computational Control, Signals, and Circuits* 1:1.
- De Jong, K. 1975. *An Analysis of the Behavior of a Class of Genetic Adaptive Systems*, Ph.D. Thesis, University of Michigan.
- Donoho, D. 1993. Nonlinear Wavelet Methods for Recovery of Signals, Densities, and Spectra from Indirect and Noisy Data, *Different Perspectives on Wavelets*, American Mathematical Society.
- Galaviz-Casas, J. and A. Kuri 1996. A Self-Adaptive Genetic Algorithm for Function Optimization, *Proceedings of the Ninth International Symposium on Artificial Intelligence*, 156-161.
- Goldberg, D. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison Wesley.
- Goldberg, D., J. Horn, and N. Nafpliotis 1995. A Niche Pareto Genetic Algorithm for Multiobjective Optimization, *Proceedings of the First IEEE Conference on Evolutionary Computation, IEEE World Congress on Computational Intelligence*, 1: 82-87, IEEE.

ISO 1994. ISO/IEC 10918-1 and ITU-T Recommendation T.81, *Information Technology – Digital Compression and Coding of Continuous-tone Still Images: Requirements and Guidelines*.

ISO 1996. ISO/IEC 14495-1 and ITU-T Recommendation T.88, *Information Technology – Lossless and Near-lossless Compression of Continuous-tone Still Images*.

Mallat, S. 1989. A Theory for Multiresolution Signal Decomposition: The Wavelet Representation, *IEEE Transactions on Pattern Recognition and Machine Intelligence*, 11(7): 674-693.

Mallat, S. 1998. *A Wavelet Tour of Signal Processing*, Academic Press.

Miller, B. and D. Goldberg 1995. Genetic Algorithms, Tournament Selection, and the Effects of Noise, *Complex Systems*, 9: 193-212.

Monsieurs, P. and E. Flerackers 2003. Reducing Population Size While Maintaining Diversity, *Genetic Programming: Proceedings of EuroGP'2003*, LNCS 2610: 142-152, Springer-Verlag.

Moore, F. 2005. A Genetic Algorithm for Optimized Reconstruction of Quantized Signals, *Proceedings, 2005 IEEE Congress on Evolutionary Computation*, IEEE (in press).

Moore, F., P. Marshall, and E. Balster 2005. Evolved Transforms for Image Reconstruction, *Proceedings, 2005 IEEE Congress on Evolutionary Computation*, IEEE (in press).

Price, K., R. Storn, and J. Lampinen 2005. *Differential Evolution: A Practical Approach to Global Optimization*, ISBN: 3-540-20950-6, Springer.

Ramsey, C. and J. Grefenstette 1993. Case-based Initialization of Genetic Algorithms, *Proceedings of the 5th Annual Conference on Genetic Algorithms (ICGA'93)*, 84-91, Morgan Kaufmann.

Taubman, D. and M. Marcellin 2002. *JPEG2000: Image Compression Fundamentals, Standards, and Practice*, Kluwer Academic Publishers.

Walker, J. 1999. *A Primer on Wavelets and their Scientific Applications*, CRC Press.

Wang, Z., A. Bovik, H. Sheikh, and E. Simoncelli 2004. Image Quality Assessment: From Error Visibility to Structural Similarity, *IEEE Transactions on Image Processing*, 13 (4): 600-612.

LIST OF ACRONYMS

AFIT: Air Force Institute of Technology, Dayton, Ohio
AFRL: Air Force Research Laboratory
D4: Daubechies-4 wavelet
DNA: deoxyribonucleic acid
DWT: (forward) discrete wavelet transform
DWT⁻¹: inverse (reverse) discrete wavelet transform
ER: error ratio
FSR: file size ratio
FS: file size
G: the maximum number of generations for a genetic algorithm run
GA: genetic algorithm
ISO: International Organization of Standardization
JPEG: Joint Photographic Experts Group
LL: vertically low-frequency and horizontally low-frequency
M: the size of the evolving population for a genetic algorithm run
MR: the number of multiresolution levels
MRA: multiresolution analysis
MSE: mean squared error
NY: New York
OH: Ohio
p_c: probability of crossover
p_m: probability of mutation
Q: quantization
Q⁻¹: dequantization
RRS: Rome Research Site, Rome, New York
SE: squared error
VFRP: Visiting Faculty Research Program
WPAFB: Wright-Patterson Air Force Base, Dayton, Ohio